TSMC's 30th Anniversary Celebration Forum
Semiconductors: the Next 10 Years

<u>Panelists' Talk</u>

TSMC Chairman Dr. Morris Chang:
First I'd like to thank all of you for joining us, coming here to join us in our 30th anniversary celebration. I particularly want to thank the panelists that we have invited for this forum.

This is a very distinguished panel, and I will introduce each of them before each of them speaks. The forum will consist of two parts: the first part is the big one, it's about two hours, and each panelist will speak for about 10-15 minutes on a topic of his own choice, as long as the topic is relevant to our theme, which is "Semiconductors: the Next 10 Years".

I have given them wide scope. They can speak on technology, they can speak on market, applications, products, on M&A, whatever, as long as it is related to "Semiconductors: the Next 10 Years".

The panelists are alphabetically arranged. Now, the first one is Jensen Huang. However, he just told me that his slides were not ready yet, so he would like to defer to a later time, and he will tell me when his slides are ready, so let's start with Steve Mollenkopf.

Steve is the CEO of Qualcomm, one of the biggest companies in the semiconductor industry. And Qualcomm, as you know, has been TSMC's partner for a long, long time. Steve, please:

Steve Mollenkopf:
Well thank you. I will try to be very brief; I think I brought one slide, and it's ready, so we'll go through it.

First, thank you for having us here and thank you for the long partnership. Incidentally, Qualcomm is 32 years old so I thought I would just talk about one topic, in a panel of this size, just leave you with one topic, which is – if you look at our first 30 years as a company, which overlapped very broadly with the first 30 years of TSMC, the main

theme was “how do we get people connected to each other?”

And in the beginning it was voice, and it rapidly moved to getting a connected computer in your pocket. And there was a tremendous amount of economic benefit that was created as a result of those things occurring, and we were obviously a part of that as were a lot of other people.

Today now, just by the nature of having a connected computer in your pocket, there are a number of new business models, companies that didn’t even exist 10 years ago, much less 30 years ago.

So for example, all we really did was we put something that had a data connection and a computer in your pocket with GPS and a camera and now you have an entire Internet ecosystem that didn’t exist before.

But that’s sort of where we are today, which is, all we did was connect people everywhere that they are. But in the next 30 years, and certainly in the next 10, there will be a tremendous opportunity as everything gets connected to the cloud, and we are obviously working on that.

There are a couple of ramifications that I think are interesting. Everyone talks about how everything is getting connected to the cloud and what will happen. From a semiconductor point of view, and certainly from a company like Qualcomm, we tend to view this as a systems problem. How do you create the underlying technology at scale to allow people to use those technologies in a way that they can focus on the business models that are available because those technologies exist?

So, sort of a similar model to what TSMC does, which is, you don’t have to worry about the semiconductor aspect of being a semiconductor company. We try to make it so that you don’t have to worry about the systems aspect of being connected to the cloud everywhere.

Now, there will be huge ramifications, I think, to the way that computing is done. And probably the one that is the most interesting and the one that I’m most looking forward to participating in is, as we get the latency down and the computing power up on devices, the ability for the cloud to become inside out and for the smarts to move closer to where the data is actually generated, which is where it’s most valuable.

I think there will be implementations and machines that will be created through the platform that TSMC enables that will look very different to what you see today. Today you can think of it as "we build central nervous systems", but in the future we're going to build nerves at the end of your finger that are going to be smarter. And what you have to engineer to do machine learning, and smarts, and a very low power process, low power machine in order to do that, will be very different.

So for us, what we're excited about is that we there are many systems opportunities or problems to solve across disciplines, and we think the big focus of our company, and I think the big focus of at least a portion of the semiconductor industry for the next 10 years is "how do we make that something that everybody in the industry can use without having to be an expert on it?"

I think it will be a very exciting time. In my mind, this is probably the most exciting time to work in the industry just because everything is now going to need connected smarts, and that's really what we do as an industry. I don't think there's ever been a time that this industry is more relevant to the big problems that will be solved in the next 10 years.

That's our one thought. I look forward to hearing from the rest of the panelists, and thank you Dr. Chang.

Chairman Chang:
Thank you, Steve. Steve's talk was much shorter that what I really wanted, but we'll wait 'til later for you to remedy that problem, OK?

Steve Mollenkopf:
I've become very good at remedies actually. I think it's a different meaning though, I think it's the second meaning.

Chairman Chang:
If you're curious about why I'm wearing this headset, I have it tuned to the English channel, it's not because I don't understand English. It's because this headset helps me to hear better. Alright, we'll give you a bye for now, Steve, OK? And the next is Vince Roche.

Vince is the CEO of Analog Devices, and Analog Devices has been one of our oldest partners. I remember that even in the middle to late 90s, we started to do pretty

serious business with Analog. Vince is actually the third-generation CEO that I know at Analog. First generation was Ray Stata, and then was Jerry Fishman, and now it's Vince. Vince, please.

Vince Roche:
Thank you Dr. Chang, great introduction. Good afternoon everybody, it's a great thrill to be here, to celebrate TSMC's 30th anniversary.

So as our name says, we're an analog company, and our business is all about bridging the world of the physical and the world of the digital. We're a 52-year old company, so we've had to morph many, many times through different types of technology.

Our business is really in the spaces we've chosen. Our business is really about scope rather than scale. Obviously a certain level of scale is important, but our business is really about scope; having lots of inherent underlying technologies that enable us to address many hundreds of applications. We actually have about 50,000 to 60,000 customers in our portfolio, and I think about 60,000 product SKUs these days as well, so a very, very complex business.

I remember not long after joining ADI in the late 80s, I remember the prognostications that the analog business was finished, it was all about digital. Our investors were confused about who we were, and whether there was a future or not, but the good news is, our business is about developing a very nuanced view of the real world in many different types of applications, developing better information, creating information and data for our customers – data about the real world, data about ourselves. So I think the analog space is a very, very critical bridge for a long time to come, and it won't be replaced, because the world will ever be physical.

So, over that period of time, as I've said, we've had to increase our scope. Actually, we started life as a module company, and then acquired some patents for data converters and become the market leader in data conversion technologies, and that brought us into many, many different types of applications.

Beyond that, our first interaction with TSMC was in digital signal processing, and that was an interesting situation for ADI because we didn't have the internal CMOS capability, and we had to look at a foundry model. But in our business, there was always a very strong sense that you had to combine the semiconductor process

development with the circuit development.

So many, many years ago Ray Stata had the vision with Dr. Chang to attempt to try this foundry model, which has been a tremendous success for us. We have followed many, many generations, many lithographies, down into 28 nanometer, and we'll certainly move beyond that, particularly in the areas of high-speed signal processing, where we got the benefit of very, very small lithographies.

I think at the present time there's a bit of confusion in our industry about "is the industry really on an innovation track, or is the industry in a slow growth period, that will never recover?" My own sense is that our customers are trying to tackle very, very complex problems. In many ways our customers are abandoning the world of hardware design, certainly the analog space, so they need the help that we bring more than ever.

Incidentally, if you look at the numbers, we produce 350-400 billion dollars' worth of revenue collectively, and sitting on top of that is the 4 trillion dollar IT sector. So all of the things that are being attempted by the cloud companies, the industrial automation companies, the car companies, it sits very firmly on that foundation of semiconductors.

So maybe over time we will introduce new modalities – we're working with carbon nanotubes at the present time ourselves, trying to build hybrid silicon-carbon nanotube technologies for specific sensing applications. But my own sense is that in spite of the prognostications about the end of Moore's Law – certainly it's got challenges, but we'll blow through those challenges, I think, through continuing to push lithographies, but also bringing new techniques to bear.

We're using heterogenous integration at the time as a company in different applications, and using more and more 3D packaging technologies to help us bring more modalities of innovation beyond the pure lithography game.

So my own sense is that the skills that we have to develop – as I said there are many, many areas where we develop 3D and heterogenous integration so we're pushing technology along the more-than-Moore curve and the Moore curve

But also if you look at the suite of applications that we're trying to solve now, more and more things are digitalizing. I think our customers are dealing with exponentially

increasing complexity. It's our job to help solve those problems for our customers as best we can.

So I think the future is more challenging, and I think one of the big challenges that we have as an industry is how to extract the value for the innovation investment that we're putting in. So I think that's an ongoing challenge for us.

But the one thing that I'm certain of is, the demand for our technology is there. The question is, what types of new techniques do we bring to solving the problems, feeding the supply, and digging into the demand available.

So I think that aside from the technology innovation, more and more we're having to find ways to match the applications and the problems we're trying to solve with the application level with the solutions we're trying to develop. And I think that's one way to deal with the cost issue, where we specialize rather than generalize.

And I think as well, on one of the TSMC charts, there was an ecosystem mapped. That's becoming a more critical feature as well of doing business in this industry. We have to get better at moving beyond the boundaries of our own companies, and developing more rich formal and informal partnerships in the ecosystem. For ADI, for example, software is becoming a problem that we have to grapple with more and more, not just the hardware.

So, those are just some ideas. I think in terms of the analog industry, the kind of innovation modes we're working along, I don't see a better solution to the problems of the world of information and communications without semiconductors.

So that's pretty much all I have to say at the present time, Dr. Chang.

Chairman Chang:
Thank you, Vince.

Next is Simon Segars, Dr. Simon Segars, the CEO of ARM. ARM is not a very old partner of TSMC's, but in recent years, you have become a very important partner of TSMC's, just as you have become a very important partner for a lot of semiconductor companies. Simon, please.

Simon Segars:

Thank you, and good afternoon everyone. It's my pleasure to be here to celebrate this incredible milestone of 30 years. As Dr. Chang said, ARM and TSMC have been partnering for a short period within that.

I think I first came to visit TSMC in about 2000. We could see an increasing number of our partners had stopped building their own fabs and they were building wafers with TSMC. And we thought about this internally and thought "well it would be a really great idea if we could partner together and learn together and trial the manufacture of our latest microprocessor on TSMC's latest process", so we started working together on that.

ARM is not quite as old as TSMC, we're about 27 years old, but I can remember the first partners we had did have their own fabs. I can remember the first time we engaged with somebody who had gone fabless. I'm thinking to myself "this is crazy. You're going to make all these chips and you don't own a fab. How are you possibly going to do that with any degree of efficiency?" But clearly it turns out it was a really good idea, one that has caught on, and we wouldn't be here today without it, let's face it.

So in coming here, I was thinking, well, 30 years ago when TSMC was formed, what was I doing? Well, I'd actually just gone to university, having gone through the home computing era and learned about technology and been excited about what computing and electronics could do together.

I was at this point in my career where I was taking for granted a lot of technology. I'd grown up in a home with color TV and a telephone, and my parents probably thought I was completely spoiled. Here we are 30 years later with all this technology that Steve was describing, and I look at my kids and I hear everyone saying "kids take this for granted, they're completely spoiled with the technology."

But you think 30 years prior to the founding of TSMC, what was going on in 1957, probably the coolest thing you could own was a vinyl record player playing 45s, and the Russians had just launched Sputnik, so they kind of started the space age. And if you think of those advances in 30-year increments, it's just incredible, it's way better than linear. We're on this exponential curve of what technology is doing for us.

When you project yourself forward 30 years, if the space age drove the '60s on, and these improvements in semiconductor process that led to improvements in computing

and increased density in memory and the ability to store everything, if that's driven the last few years, what's going to drive the next decade, the next 30 years? And I think that's going to be data.

We're going to see all these tiny devices, based on all these transistors that TSMC is going to manufacture for us, deployed anywhere and everywhere. We're going to see tiny sensor gathering data from anywhere and everywhere. We're going to process it at the edge, we're going to process it in the network, we're going to process it in the cloud, and that's going to create what I like to think of as a globally distributed mega-computer, that just going to absorb all this data. It's going to be fed by the world's data.

And even today as people are talking about big data, I think we're really just at a starting point. We're at a starting point of huge data, and quite what we do with it is yet to be determined. I think we could probably talk all day about the opportunities that come from that, and I think they are endless, and they're going to create enormous opportunities for new businesses, new business models, and will be the driver of what all of us in this room spend the rest of our careers really worrying about: How do we leverage all of this data, how do we leverage all of these technologies that have come about and make the most of it?

So I see no end of opportunity there, no end for the deployment of more and more transistors in more and more places. But one thing that could potentially get in the way of that, and one thing that I think that as an industry we need to collaborate around are two issues, and that is security and privacy.

The security issues you see raised in the press every day right now are about stealing credit card numbers and social security numbers. I think that's going to pale in significance compared to the kind of data that could get stolen and misused in the future.

So as we go into this future where data is driving the technology, we've really got to fundamentally re-think how we secure the data, really engage in grown-up discussions about the privacy models around that data because ultimately what we're trying to do is make technologies that make peoples' lives better.

We often talk about "users". Well, those users are actually people, and most people, it turns out, don't have degrees in electronic engineering and computer science, and

aren't really equipped to deal with today's security problems, let alone those that are coming down the pipe. So I think as an industry we've got to really think about this, think about how we create a new world that is a lot safer for the citizens that will live within this digital world. We've got to throw away these old models of just making something secure, testing it, and when it leaves the factory, kind of having no further obligation to look after the person using that device. I think that has to change.

You can only know that you're secure against the threats that you knew about when you finished testing the product. The problem is that security is a live thing, security threats are a live thing, and unless we engage in owning the device after it's been shipped, then I think the users, the people who use our products, are going to be at risk. And that's ultimately going to constrain all the benefits that we could have from AI. So this technology is about people, and this data should, and can, bring a lot of benefits to those people.

Now since we're celebrating 30 years, I thought I would go and talk to my daughter, and the reason for that is that she's almost exactly 30 years younger than me. And I said to her the other day "Ok, what do you want out of technology in the next 30 years?" This is somebody who I think is surgically attached to her smartphone. The only time she puts it down is when she's broken it, and I have to insert a new one in her hands. She has it all the time. I said to her "Well, what do you want?" And she said "Well, millennials don't really want flying cars and all this stuff that you old people talk about" – she was looking at me – "What we want is technology that makes everyday life just better."

And I thought, OK, good point. Flying cars and gadgets, we're kind of on that. But really thinking about how technology improves peoples' lives, which is what it's there for, is what we ought to be focusing on. So I hope through yet more collaboration, building out this ecosystem even further, we can get focused on those things, and make this technology world a safe place for everybody to go and occupy.

Thanks a lot.

Chairman Chang:
Thank you, Simon. I was informed just a few minutes ago that Jensen's slides are now ready. So let's get Jensen back in line now. Jensen, are you ready to talk now?

Jensen Huang:

I was really expecting a long and gracious introduction

Chairman Chang:
Excuse me. I haven't really introduced you yet. Jensen Huang is the CEO of NVIDIA, really a very old friend and partner of TSMC's.

I remember that back in – I think it was the middle 90s, it could have been '93, '94 – TSMC was already progressing, we were on our way to being a billion-dollar sales company. In '93-94 we were maybe a few hundred million dollars. I got a letter in Hsinchu from a person who called himself Jensen Huang, and he said that "I'm interested in talking to you about foundry business, so will you please give me a call."

So the next time that I went to the U.S., I called Jensen, and on the phone I heard Jensen saying on the other end – there was a lot of noise, you know. Apparently there was a group of people talking there. After introducing myself, I heard Jensen say "Hush! Quiet! Morris Chang is calling me!" And that started our business. That started our partnership, and it has been more than 25 years now.

Now NVIDIA, recently the stock has done awfully well, it has gone up 5 or 6 times in one year, and it's supposed to be very good at AI. So Jensen, tell us more about it, or anything else.

Jensen Huang:
Morris, first of all, the reason why I had to write you a letter was because I called your U.S. sales office repeatedly with no success, because your sales people were out calling on real customers.

Chairman Chang:
Well, I'm very glad that the current U.S. manager was not there at that time.

Jensen Huang:
And just to put it on record, the reason why my slides were late was apparently, I was the only person that prepared. (To other panelists) You know, our performance today is graded.

After all these years, like all of you in the audience, we all work super hard every time that Morris asks to meet to impress him, so hopefully today I'll say something smart.

And also, when Morris and I first met, it was 1993-1994, I was in my early 30s. I was talking to Morris today, and he said that he was only several years older than I am now, then. But we got along just fine, he's got a great sense of humor.

So let me talk to you about what I'm going to talk to you about. You know, the computer industry has gone through a lot of changes over the years. And in each one of the eras of computing, there were some fundamental forces at work. The fundamental forces that were at work typically appeared on two sides of the computing stack. On the one hand, how is computation done, and on the other extreme, how is software developed. Usually, one of those two ends of computing would be changing at a time. If you go back and think about how PC and mobile and the cloud era have evolved, that is surely true.

Now for the first time in the computer industry, and I've been in the computer industry for quite a long time now, for the very first time, two fundamental forces, one in software and one in computing, are simultaneously governing the future of this industry.

The first is the discovery of the new form of computing and new form of developing software called deep learning. Deep learning is, on its surface, rather simple to understand. It's otherwise called an artificial neural network. Basically, a large system of multivariable, high-dimensionality, linear equations. And it has millions and millions, hundreds of millions of parameters. This equation is extraordinarily large. And it's high-dimensionality, it's more than three dimensions, more than four, more than five, it's extremely high-dimensionality.

It has the ability to curve it against any type of equation, any type of condition that you can imagine. And its magic is that it can predict the outcome – classify, predict, recognize insight from very, very complex unstructured data.

For example, if you were to take a picture of this room, what it looks like in binary, in ones and zeroes, is just a bunch of goobly-glop, it just looks like noise. But from that noise, deep learning can learn how to recognize insight.

Deep learning has this ability to automatically write software, essentially, so software is writing software, and for the first time, computers can write software instead of computer scientists writing software.

Its one enormous handicap, and the reason why it's taken so long to become popular, is because it requires an enormous amount of data to learn from so that it can recognize the patterns that are existent in the data. And because it requires an enormous amount of data to learn from, and because the equation is so large, this linear system of equations is so large that it requires an enormous amount of computation in order to effectively learn how to predict outcomes from the data.

The second dynamic that's happening at basically the same time, as it turns out unfortunately, CPU performance is no longer scaling. CPU performance has stalled, what some people describe as the end of Moore's Law. Even though the number of transistors we're getting from each and every node continues to grow, the problem is that we no longer know how to take those transistors and transform them into effective computational performance for CPUs. We've essentially run out of ideas.

After all of these years, we've tried very deep pipelines, we've tried multi-scalar, basically multi-issuing of instructions, we've tried speculative execution, out-of-order execution, we've tried basically all the techniques we know how, and we've wrung out of the CPU as much performance as we know. As you see, the white line is transistor progress – as you shrink geometry, the number of transistors continue to grow, and on the blue line is basically CPU performance.

Now, what's really staggering about this is that CPUs are basically improving by a factor of 5 every 5 years. So basically every 15 years, it increases in performance by a factor of over a hundred. If we lose that trend, we would have a deficit of 100X in another 10 years. And so 100X in performance at a time when the world has discovered this incredible new software technique would be a real tragedy.

These two are essentially the governing dynamics of the computer industry today. Almost every software company in the world is trying hard to adopt deep learning because of its incredible effectiveness, and I'll show you some of its results. And then secondarily, we all recognize now just riding on the CPU's natural performance curve is no longer an effective way to continue to improve performance.

The reason why NVIDA has been growing so fast is because we pioneered this approach called accelerated computing some time ago. This is almost 15 years ago now. The fundamental business of NVIDIA is about creating a computer architecture that accelerates applications that otherwise wouldn't be possible. And the way we went about doing that was adding a processor as a companion processor to the CPU.

Instead of replacing the CPU, we had the wisdom of offloading the CPU of the work that the CPU wasn't good at doing, where we could do a much, much better job.

The architecture of our GPU is fundamentally different from the CPU. The CPU is a sequential processor, we're a parallel processor inherently. The performance results, the benefits, have been extraordinary. And in fact, just this year, two Nobel Prizes – one is called cryogenic electron microscopy, basically using electron beams to image biomolecular molecules, and then the other one is called the laser interferometer gravity wave observatory; basically two massive arms, two tunnels that light all the way down the tunnel reflects back at a mirror, and they were able to prove something that Einstein predicted 100 years ago, that gravity had waves. They detected two black holes colliding in 2015, and they were just awarded the physics Nobel Prize.

In both of those examples, the GPU was used as the computational engine for these incredibly high-fidelity imagers, and without the GPU it surely wouldn't have been possible. So it's great to be able to contribute to both of these Nobel Prizes.

Now let me show you how other people are using it. In industry after industry, whether it's in high-performance computing, in life sciences, in internet services, our GPUs are now used in everything from voice detection, speech, when you converse with your phone, the GPUs are used to train the neural networks used to recognize your voice, to high-performance computing, to discovering new drugs.

But let me show you a few examples, I made a couple of home movies for you. In this case, the first one I'm going to show you – and Simon, I've got to talk to your daughter; I think flying cars are way cool. As you know, we finally figured out that one of the most important technologies to enable self-driving cars. In order to create self-driving cars, the fundamental problem is perception. The car has to be able to recognize what's happening around the car, and determine what to do: first of all, perceive the surroundings, second, reason about where it is relative to the world, and third, make a decision about how to drive. Perception, reasoning, and planning are the three fundamental functions and actions of robots, artificial intelligence agents. And for the very first time we've been able to solve those problems because of deep learning and with GPU acceleration. I'm going to show you some of these results. This is a homemade movie that gives you a sense of what's possible.

[Self-Driving Car Video Plays]

The processor is running this neural network and the neural network looks at all of these pixels on the screen, and automatically recognizes what's a truck, what's a car.

The reason why they have segments – the distance from the car.

Just imagine, the car is looking at this video. These are the neurons that are inside the brain of the artificial network.

This is the car seeing everything around it, surround perception. It doesn't matter what the orientation of the car is, it recognizes it as you drive by it. It recognizes people, signs, lights.

Imagine, what the car ultimately sees is a bunch of ones and zeroes. It recognizes cars, but it also recognizes where it's safe to drive. It's called safe space.

It processes LIDAR, which is a bunch of point clouds. It figures out where it's at, it's called egomotion, visual odometry. It figures out where it's at relative to an HD map, compares itself to an HD map, uses the LIDAR to figure out where it's at.

So it's perceived the surroundings, it's figured out where it's at in the map, now it's got to choose its path. It's got a bunch of choices to choose from, and it chooses the safest, the best ones.

It comes up to cross traffic, looks side by side, and makes a hard left.

Artificial intelligence – self-driving cars.

[Self-Driving Car video ends]

Ok, I'm going to show you one more example of what artificial intelligence can do.

Artificial neural networks are at some level, an extremely complex concept, at some level, relatively easy. The question now is, how do we turn this incredible new technology called deep learning into something that ultimately, reflects our ultimate version of intelligence, which is autonomous characters.

The example I'm going to show you here is called Isaac. You're looking at a screenshot of the world that Isaac is currently living in. And here what we're going to

do is try to teach Isaac how to putt. The way to think about putting is, you've got to pick up the club, you've got to figure out how to strike the ball, where to strike the ball, you have to read the green, because as you know, the green has slope. And you have to figure out how to hit the ball, eventually landing it in the cup.

So the question is, what's the methodology for doing that? We're using a technology here called deep reinforcement learning, basically the same way we learn as children. We try, try, try, try, try again, and every time we fail, we're punished, and every time we succeed on a random basis, we get rewarded, and over time we work towards the goal.

OK, so I'm going to show you a home movie here of Isaac learning how to putt.

[Machine learning video plays]

There's a slope on the green.

This is how Isaac started, just trying to see "where is the ball? Where's the cup? How do I swing this club?"

And he accidentally strikes it – that's good. Now do it the same way next time.

And now it's figured it out.

Look at this – the ball's obeying the laws of gravity.

Ladies and Gentlemen, artificial intelligence.

[Machine learning video ends]

Of course, what you're seeing there is a virtual world that obeys the laws of physics where Isaac can see the world and it's completely photorealistic. And the goal is, eventually we will take the trained brain of Isaac and put it in a real robot, put him on a real green, and technically, theoretically speaking, Isaac should be able to putt.

This looks a little bit like a toy right now, but as you know, the beginning of the personal computer industry started out as a toy, and many of these revolutions, the earliest examples of what's possible, is toy-like. But if you can extrapolate into the

future and imagine how good it can be, once we can solve this problem, obviously manufacturing robots will be highly automated, surgical robots will be highly automated, the way we farm, the way we do construction, the things that are dangerous to people can now be replaced by robots.

It's very clear now that this is the era of AI. I'm sure all of you have been hearing about the power of AI in industry after industry, investing in AI. At the highest level, the reason why I think it's so incredible is this: if farming is the automation of food, and the industrial revolution was the automation of power, AI is the automation of automation. And this is the era where computers can learn by itself, software can develop software by itself, the potential is really quite amazing. And as I've showed you just now, the two examples that I've demonstrated are actually software that no human knows how to write. To imaging that there's software that no human can write, but a computer can, is really quite a truly remarkable thing. And so I think that these are exciting times.

I made a birthday card for you, Morris -- To Morris Chang and TSMC: In celebration of TSMC's 30th birthday, Morris Chang's extraordinary career, and our lifelong partnership, happy birthday from all of us.

Chairman Chang:
Thank you

Jensen Huang:
One of the most touching things that anybody has ever said to me, Morris, was a researcher who came up to me and said "Jensen, because of your work, I can do my life's work in my lifetime" because of how we accelerated quantum chemistry for him.

And I can say with complete certainty that because of your work, we can do our life's work, so thank you very much.

Chairman Chang:
Thank you, Jensen.

Now we get back to the line again. The next is Hock Tan. Hock Tan, the CEO of Broadcom. Actually our partnership with Hock Tan's organization dates back to before Broadcom. It dates back to the old Broadcom, of course, Avago, that Hock Tan was the CEO of, and even before Avago – LSI Logic, and the company that went back to…

Hock Tan:
Integrated Circuit Systems

Chairman Chang:
Oh yes. Now, Hock and I also share something else in common, that is that he and I are both Bachelor and Masters of Mechanical Engineering from MIT: same major, same degrees, same school.

Hock Tan:
Different times.

Chairman Chang:
(Laughs) Different times. He told me this morning that he was class of ’75, and I was class of ’52. OK, Hock, please.

Hock Tan:
Thank you Mr. Chairman. Distinguished guests, in the presence of this technological insight that I’ve been imbibing, it’s with some trepidation that I’m offering to share some thoughts on where our industry, the semiconductor industry, is headed.

By the way, for logistics purposes, I’m doing a series of slides, so you’re better off twisting your arm muscles to look at the screens than look at me. So let’s start.

As with most things, the story begins with demand. Looking back over 40 years’ span, you see an industry that grows in the 80s, 90s at over 30% every year. And yet today it grows at less than 5% a year, basically GDP rate.

And if you look at the slide below, as in any gold rush, the semiconductor industry delivered to early investors’ alpha returns – very high returns, above the S&P 500. However, the last 10 years the returns have regressed to mean performance. In other words, the gold rush, as we know it, is over.

Structurally, this industry has evolved very interestingly. In the 80s, it was start-ups, small start-ups. In the 90s, we have IPOs and small tucked-in strategic mergers and acquisitions. Then in the 2000s there was a frenzy of carve-outs from large OEMs who had built up semiconductor capability within them. So there was a frenzy of carve-outs, leveraged buyouts, which kind of obscured the fact that the industry was slowing down. And for the last 5, maybe 10 years, we’ve seen those horizontal

mega-deals. Consolidation, as I call it.

But just because the gold rush is over does not mean we cannot make a lot of money. Case in point – I can't resist this – Broadcom. That's the Broadcom revenue trajectory over the last 10 years, and if you look closely at the last five years, I'll be honest, there's some growth, but we struggled mightily in this very competitive hostile environment to just even eke out that level of growth.

Then, in 2013, the light bulb turns on, so to speak. Basically, we went out there and started acquiring all those great businesses that we see out there in semiconductors. And there are lots of them – it's a target-rich environment. It also would not have been possible without the support, help, leverage on the strong enhanced partnership with TSMC, and you can see the last 5 years as history.

As far as TSMC is concerned, they didn't do too badly too. It was a very great partnership. But this a walk down memory lane, I know, to you guys, and no walk down memory lane is complete without – let's take a look at the ecosystem, the semiconductor and technology ecosystem as we see it.

So, you see there as an inverted pyramid, pre-1987. I picked '87 for a good reason, you'll probably figure it out sooner or later. It was the dark ages. Huge OEMs ruled the earth. It supported fragmented demand from enterprises and consumers.

Then in the 90s, light shines. Driven by mega-trends from PCs, and later on mobility, the semiconductor industry grew, and grew very fast. Innovation was rife, and that innovation drove an interesting phenomenon, which is foundries. The emergence of foundries recognized the distinction between product design and intellectual property, and process, and starts the emergence of foundries and semiconductor growth.

Then in the 2000s, data centers started happening, and that continued to drive semiconductor growth, albeit at a reduced rate. What's also interesting is foundries benefitted from that era of OEMs shedding the semiconductor operations of theirs.

And finally, the last five years also, you basically see in the industry semiconductor demand slows down to GDP levels, and horizontal consolidation occurs. The mega-deals happen, as I said. Consolidation happens. Supply, I hate to say, stabilizes. Foundries get institutionalized. So we got ourselves, at optimum, some basic level of stability.

But if you look at demand, things are changing. You see social media driving growth of cloud. You see demand of enterprises manifesting through cloud, so today you have an ecosystem that sits where spending on cloud is virtually half of total IT spending up there.

So what's next? And here's how I see it – very likely scenario, vertical integration. The cloud operators have reached unprecedented scale. They have the scale to reach down vertically in our ecosystem stack. They have talent to do systems architecture, they write their own operating software, do they need OEMs? More than that, they can do the chips. They have the capability to draw talented engineers, design engineers in silicon. Do they need the design IP sector?

Now, going down one step further in the stack to foundries, or to silicon production – the financial requirements might give them some pause. But I doubt it unless we see the foundries reaching out vertically.

Thank you.

Chairman Chang:
Thank you Hock. Next is Peter Wennink, the CEO of ASML. Actually, ASML shares quite a few things in common with TSMC. First, we're about the same age. I think ASML started one or two years before we started.

The second commonality is that Philips was the major investor in ASML when it started, and Philips was the second-largest investor in us when we started. And our partnership dates way back to the time we started. Actually, Philips suggested to us that we buy from ASML, and ASML was of course happy to accommodate it. And ever since we started, I think ASML was almost the sole source of our lithographic tools for all these 30 years. Now Peter, Please.

Peter Winnink:
Thank you Morris. Now, when you asked me to say a few words, I was humbled, because we're a supplier, and we do our best. But indeed, when I look back and see what constitutes the relationship between TSMC and ASML, it is partnership. You mentioned this, and I think when you asked me to look forward the next 10 years, of course I'm going to say something about scaling, but I think it's increasingly important that we realize that we need a level of partnership for innovation, because things are

not getting any easier, and this is what I want to talk about.

I think our industry is very well-positioned to tackle some of these big challenges that we see in front of us. And everybody knows it, whether it’s sustainable energy, or it’s the food and water supply, or the aging population, or the massive urbanization that we will see coming at us. Our industry is very well-positioned to do this.

And have we been successful in finding solutions? I think we have. And have we been able to create value? I think we have. Looking back a bit, if you look at this industry, this industry constitutes of semiconductor equipment manufacturers and then semiconductor design companies and the manufacturing companies, going to hardware and then services. If you look at the entire ecosystem, the entire ecosystem has an earnings potential of, this year, close to 300 billion dollars of EBIT. That’s massive, and it means that that there is such an enormous amount of value that is out there, and the research and development capabilities are massive. So we’re very well-positioned to take advantage of our position, but also help these big societal challenges.

What has driven this industry? Scaling. Scaling has actually helped us a lot, and Jensen showed a chart that also showed it very much going up together with the CPU power, so scaling has helped us tremendously. We’re currently able to increase transistor density, we’re at levels where we’re at over 100 million transistors per square millimeter, and that has been driven by scaling.

When we think about scaling, and think about ASML, you basically say there’s one level of scaling; it’s geometrical scaling, basically shrinking through pattern and design optimization, but that’s not the way we see the world. There are different types of scaling that actually help us get to the point where we want to be. It’s device scaling, and in fact if you want to make this a bit more simple, device scaling went from planar to finFETs and to nanowires, and from a geometrical scaling point of view we designed lithographic equipment which went from i-line to ultimately EUV. I’ll talk about that a little bit later, but we also need the circuit scaling. We went from individual ICs to a very wide variety of chip types. And the architecture scaling is just as important. It went from general purpose processors to more dedicated silicon.

But all these four things, they do not happen independently. They have to be combined. It actually means we need to work together very closely with our customers to understand the impact of these different scaling types on what we’re doing.

And even in geometrical scaling, we need the help and support of our industry partners that are doing deposition and etch and all the other steps. Sometimes when you talk to the financial world they try to polarize who gets what, who wins and who loses, which basically is ridiculous. We basically need all of it. We need to partner more strongly than we did before because things are getting more complex.

Now when you're talking about complexity and looking through the rear-view mirror, you think that everything was easy when we look back. Well, it wasn't the case. When we look back from where we are today, we went through big challenges also. And some of those challenges at that time, looked like they were very difficult to overcome, but we did overcome them, and we did it by working together.

I couldn't resist looking back; when Morris asked me "why don't you talk about the future?", I said I also want to look back in history because history tells you a lot about what the future might potentially bring.

The first stepper – you mentioned it Morris – 28 years ago, we shipped the first machine to TSMC. And I asked our people "where is that machine?" because I know that 97% of our machines are still working. Now, this particular machine, that first that we shipped to TSMC, is still here although it donated its parts to other machines of that same generation. But the others are still working. I'm just saying this because over these 28 years we've done some incredible things, and I think it was due to this partnership.

Now when we now look at today, we've moved from 0.8 micron now to single nanometers. We went from g-line to EUV. And EUV, as we all know, has been a rough ride until a couple years ago and we've made a lot of progress.

So, what will EUV bring us? Well, EUV brings us, in the foreseeable future, a billion transistors per square millimeter. And if we think about 2030 and what can next-generation EUV do, we can even go beyond that.

And then if we look forward over the next 10 years, and see what are the challenges we have overcome, you might remember the first immersion tools that we shipped to TSMC, which basically were tools that didn't perform that well when we had the first tools out there, but with the help of TSMC and other partners we made it a very reliable tool and it is, in fact, the workhorse of the industry today, and has enabled us

to do multiple patterning that we needed as a bridge solution because EUV was not really ready. And we could have only done that because there was such a close cooperation and sharing.

Now we're at the age of extreme ultraviolet, EUV, and we're smack in the middle of the preparation for volume production. And it's going to happen, because EUV is going to add simplification, shorten R&D cycle time, it will reduce work-in-process. In other words, it will do what Moore's Law is, which as an empirical law of economics, will reduce cost.

We've started last year and the year before into the basic research into the next level of EUV, High-NA, and it will happen. We have the optical designs, we know what to do, we will use the same EUV source. High-NA is there to help us in the second half of the next decade, it will bring us beyond 2030. So from a scaling point of view, we have the instruments, we have the equipment, but what we need is our partnership.

You know, there are always bumps in the roads, and that's why I said earlier when we look back in the rear-view mirror everything looks fine, but at that time we did have our issues. And we overcame them. Just like it was a long road to make EUV a production-worthy tool, and especially the last couple of years we've made huge, huge, progress. And that has been the case because we have the partnership.

I want to highlight four of the basic components of the partnership that we have with TSMC. This we would like to extend, and are extending to other industry partners also, and it's all about building partnership based on trust. That's not an old-fashioned thing. Trust is real, and that's really the four pillars of the partnership we have with TSMC.

It is Competence. We employ the brightest engineers, we make them work together, and we can do something very, very special.

Secondly it is reliability. As a partner you need to be very reliable. It means that we are going to tell each other what we're going to do, and we're going to do what we tell each other.

Thirdly, and I think it's very important, is transparency. When we have to deal with those difficulties and those big challenges, it's extremely important that we are transparent with each other so we understand what needs to be done, what needs to be done now but also in the next-generation machines in 5 years' time and in 10 years'

time. And that is not always easy, to be transparent, but it's absolutely essential if you want to overcome these big technical issues.

And the last and probably most important, we have to manage our self-interest. Especially in the relationship with TSMC, it's about risk and reward sharing.

Those four building blocks form the foundation for us making sure that we can overcome those big, big challenges, and I'm convinced that where we are today with EUV, with the partnerships that we have, we have a very bright future in the next 10 years, and I think together we can make it happen. Thank you, Morris.

Chairman Chang:
Thank you, Peter. Now last, but certainly not the least, far from being the least, is Jeff Williams, Chief Operating Officer of Apple. Our partnership with Apple did not go back very far, but it has been intense, and it's important. Certainly very important for TSMC, and I've been told that it's very important for Apple too. Jeff?

Jeff Williams:
Thank you Dr. Chang. I was going to run through our very detailed product roadmap for the next few years and ask all of you to just keep it amongst us, but like Jensen, I had some slide issues, so I'm just going to ramble on for a few minutes instead.

First off, thank you, it's a real honor to be here with this distinguished group, and we're here of course to celebrate TSMC's 30 years. It's amazing, as you've seen in the slides, how far technology has been driven over that time TSMC got its start shortly after the introduction of the legendary Cray 2 supercomputer, and 25 years later, we put the same processing power into peoples' pockets with an iPhone 4 in 2010. It really is remarkable.

It was actually 2010 that the first seeds of our partnership between Apple and TSMC were planted. I had flown to Taiwan and had dinner with Dr. Chang and Sophie at their house, and it was a wonderful dinner. We were not doing business with TSMC at that time, but, we had a great conversation. We talked about the possibilities of doing stuff together, and we knew the possibilities were great if we could take leading-edge technology and marry it with our ambitions.

What seems obvious right now, wasn't then, because the risk was very substantial.

The nature of the way Apple does business is we put all of our energy into our new products, and then we launch them. If we were to bet heavily on TSMC, there would be no backup plan. You cannot double plan the kind of volumes that we do. We want leading-edge technology, but we want it at established technology volumes, so that may be what Dr. Chang is referring to when he mentions 'intense.'

For the TSMC side, it means a huge capital investment. It means ramping faster than the more careful yield plan that the industry is used to. But together, we decided to take the bet, to take the lead, and Apple – this was our first engagement — decided to have 100 percent of our new iPhone and new iPad chips — application processors — sourced at TSMC.

TSMC invested 9 billion dollars and had 6,000 people working round the clock to bring up the Tainan fab in a record 11 months, and in the end, the execution was flawless. We've gone on together to ship over half a billion chips together in that short window, and I think TSMC has invested 25 billion — 9 billion dollars on that very first venture — there are very few companies in the world that would spend 9 billion dollars in capital across everything and not a single debt.

So for that, we thank you, Dr. Chang, and everybody at TSMC. It's been a wonderful partnership.

Now the request was for me to describe the next 10 years in silicon, which I'm completely incapable of doing. So I'm going to reframe the question. I think in politics they call this a "pivot".

It's interesting, when we look back a decade ago, the question we had was 'do we have enough processing power in our silicon to match our ambitions?'

The big challenge we had as we moved into the mobile revolution was this tradeoff between performance and power. The view at the time was you had to choose. You've got one or the other.

Largely as a result of what the fabless model has done, what TSMC has done, what many people in this room have done, Simon and his organization from ARM, we reached the point where those tradeoffs are not necessary. We have performance in thermally constrained environments. This opens up for the next decade a whole new world.

So for the next decade, the question is not so much “do we have enough processing power to meet our ambitions” – though we need to keep working, we need to have better lithography, don’t slow down. But I think the question for us is “do we have the right ambitions to go utilize this technology in front of us”.

We at Apple are not concerned about the talk of the slowing semiconductor industry. That’s not the case at all. We think the potential is huge. We believe strongly in both the cloud side, but the future will be a lot of on-device processing. We believe this is the best way to deliver great features without sacrificing responsiveness, privacy and security.

We see in our brand-new A11 bionic chip, which is made right here at TSMC, every time somebody takes a photo, there’s over a hundred billion operations. That’s just mind boggling – a single photo, over a hundred billion operations.

The potential is limitless. We’ve put a neural engine on the chip, and I won’t repeat some of the things that Jensen shared, we have the same view and vision of the potential of AI to deliver a much safer and efficient autonomous system. The neural engine on our chip has already enabled face ID processed locally.

We view that the next 10 years is about the ambition to do what Simon’s daughter is asking for, to make life better. Probably one of the most significant examples of this is our opportunity to use transistor technology advancements and power scaling to revolutionize healthcare.

We think the industry is ripe for change. We think there is tremendous potential to do on-device computing, to do cloud computing as well, to take that learning and through deep learning, machine learning and ultimately artificial intelligence to change the way healthcare is delivered. We can’t think of anything more significant than this.

I think the question in front of us is “do we have the right ambitions, and can we go do this?” There is no such thing as autonomous innovation. Human beings dream it, human beings drive it, sure we’ll have deep learning, but there’s not autonomous innovation, so it’s up to us, this generation and the next 10 years, to take advantage of what is in front of us in the silicon world.

We at Apple are really inspired. Those of us that started many years ago and sat in

front of a green monochrome computer screen we're super inspired with the state we're in. I'll just say this: if in the next 10 years, from a society standpoint, we do just a few 'gee whiz' things like flying-car kind of dreams and then the rest of the time, we're using faster chips to do the same things we are doing faster, we will have squandered one of the biggest opportunities in front of us.

I think we are at an inflection point. Much like my colleagues, with on-device computing coupled with the potential of AI to really, really change the world. We couldn't be more excited about it at Apple, and thank you for your time.

<u>Q&A Session</u>

Chairman Dr. Morris Chang:
So we resume our panel discussion. First I want to ask the panelists – do you have questions for each other?

No questions for each other. Actually, I have a few questions for you, but I look at this pile of audience questions, and there are enough of them to keep all of us busy the rest of the session, so I will just start with the audience's questions.

There are quite a few questions about AI. The first one is "When will AI be ready to predict the next 10 years in semiconductors?" Well, Jensen, do you want to take a stab at that?

Jensen Huang:
Everything I said today was predicted by and AI. Artificially intelligent, that's me.

Chairman Chang:
So you haven't given an answer.

Jensen Huang:
There are two types of AIs, and there are surely more categories than that. But an easy frame of mind is that there are specialized AIs, AIs that are designed to do one specific skill, for example, the ability to recognize images, the ability to recognize speech, translate speech, have a conversation with you, be able to predict stock price.

Those kinds of AIs are highly specialized, and I think that over the course of the next 10 years, we're going to find that specialized AIs, AI radiologists that are reading CT scans…over the next 10 years, specialized AIs will surely achieve superhuman levels.

But the other form of AI, which is general AI, I think first of all we don't even really understand how intelligence really works at the general level. General AI, or what is otherwise known as artificial general intelligence, there's no evidence that anybody has an understanding of how to get there anytime soon.

Chairman Chang:
Actually, I can answer the question. Nobody asked me, but I will answer the question. On this point, the next 10 years in semiconductors, I think the questioner is asking the market for the next 10 years. I don't quite agree with – I think it was Hock – that said that it was just GDP kind of growth.

I think that the semiconductor market will grow 200-300 basis points more than the world GDP. That's my answer.

Keep in mind world GDP is now only 2.5-3%. In a good year, it gets up to 3%. Not a good year, like this year is not a particularly good year, 2.5%. And my prediction for the semiconductor market in the next 10 years is that it will grow at 4.5-5.5% compounded annual rate.

Jeff Williams:
C.C., that's your performance plan. It just got set.

Chairman Chang:
Our performance plan is far better than that. Our performance plan, incidentally, is 5-10% annual growth.

Next question from the audience, again on AI. This one is from Bill Lu, who's an analyst from UBS. "How does the trend for AI chips play out? Will it stay discrete, or does it get integrated into an SoC like GPU, DSP?" Bill Lu, you are here, right? Why don't you stand up and ask the question yourself?

Bill Lu:
Hi, I'm sorry about the grammar and handwriting, but the question is fairly simple. If you look at AI chips versus what we've seen previously, for example, in the PC we had the math co-processor at one point and eventually it got integrated into the CPU. If you look at the apps processor, it's a lot more than a CPU. It integrates the GPU, it integrates the DSP and other things. So, longer-term, do we think that the AI processor stays discreet or does it get integrated? Why or why not? Thanks.

Chairman Chang:
You want to take a crack at that also, Jensen?

Jensen Huang:
Sure. First, Bill it takes a lot of courage to stand up after that question formulated that way. I would have sat down.

There are three markets from a chip perspective. There's the first market, which is developing the AI network. And there, basically what you're doing is, the deep learning network is trying to guess what the right answer is. And every time it guesses wrong, there is a process that computes the error back to the beginning, so it changes all the neurons, it changes all the synapses if you will – they call them weights and activation layers. And it would do this trillions and trillions and trillions and trillions of times, until eventually it's repeated itself so many times the weights and activation layers are able to guess the output. You show it a picture of Morris, and it says Morris.

It's so diverse that you can show it a picture of Morris when he was 55, and now it will still guess that it's Morris. So this network is very robust, it's very diverse, and it takes a long time to train.

In that particular situation, the number of transistors, for example, that's in a GPU for training – call it 100 billion transistors to every 6 billion CPU transistors. So it's not sensible to integrate 100 billion transistors and all of the memory that it's connected to because it needs so much bandwidth into the CPU.

In that particular case, I think for training systems I think discrete is the right answer.

Then there's the second category of neural net processing, and I would say that is the category of autonomous machines. Once you've developed a neural network, you apply it into a machine, and that machine could be a self-driving car, that machine could be an autonomous machine like an iPhone, that machine could be a robot.

In those particular cases, the neural network processing is quite intensive, and so I think that you'll find that there are some that are integrated and some that are discrete. But it's very clear that autonomous machines will see a new type of processor.

We're developing a new processor. It's called Xavier; we're in the process of taping it out to TSMC right now, and it's a new class of processor that takes real-time sensory input of all different types, fuses them all together, and it does parallel computing, CPU processing as well as neural processing all at the same time. In those cases, it could be discrete, it could be integrated.

And then the third class, I would characterize as IOTs. These are Fitbits, these are thermostats, pressure sensors, vibration sensors. They're sensors of all kinds, and I would think that Analog Devices, Vince, you guys have tons and tons of these sensors, and they're going to have these neural network processors connected to them, and they're all going to be completely integrated.

There's going to be a trillion of these sensors; Vince is going to sell these things like dust, it's going to be literally everywhere – every vending machine, every elevator, every coffee pot, every watch, every tennis shoe – you know, everything that you wear will have a little tiny neural network in it and in that case it will be integrated.

There will be hundreds of millions of autonomous machines, there'll be tens of millions of training systems and inferencing systems, and there'll be trillions of these IOT devices. So that's kind of the way you have to think about it, depending on application.

Chairman Chang:
Anyone else want to add to what Jensen said?

Vincent Roche:
Maybe if I could have the real-world in view, the physical world into the digital world view, our sense is that artificial intelligence will be pervasive, at the cloud and the GPU/CPU, but also many of the things we're starting to plan and do now at the sensor level, as Jensen said, will have some level of artificial intelligence or machine learning right down at the sensor itself, so the partitioning of AI, from our perspective, will actually start at the sensor all the way to cloud.

Chairman Chang:
Steve?

Steve Mollenkopf:
Well, same view. I think if you look at it as a continuum from the cloud here, very, very specialized, a huge amount of processing, moving to the edge of the Internet, it tends to be more integrated in that direction.

But we tend to think of it – at least for us there's going to be a lot more computing near where the data actually is, so even closer to the sensor, and I think that that's going to be a big opportunity for the people in this room, actually.

Chairman Chang:
Simon, anything?

Simon Segars:
Yes. I think as Jensen said, there are going to be a very wide range of applications. There are a ton of algorithms, there are a ton of algorithms yet to be discovered, tons of ways to utilize machine learning algorithms. Some of those are going to need small chips with a bit of acceleration, some will just run on a CPU, as bad as they are at doing these things.

And at the other end of the spectrum you've got enormous dedicated processors that are there for training the networks themselves. So I think the more than industry, scientists, and engineers get to play around with this technology the more use cases there will be and the broader range of implementations of machine learning and AI we will see. Some will be

integrated, some will be stand-alone.

Chairman Chang:
Thank you. The next question is also about AI, but it's of a different nature. It's from Paul Shrock. Paul is a consultant. Actually, he works for us. At least partially, he works for us.

He helps us to assess our customers' opinion of us every year. In fact, just this weekend I finished reading, or at least leafing through two big volumes that he just submitted to me. I read his reports every year, in addition to listening to his verbal reports.

Every time I read those reports, I can't help getting the feeling like the Irish. No matter how hard you try, the World does not like you.

Anyway, Paul Shrock is in the audience, and his question is: What responsibility do we, as technology providers, have to ensure that AI and deep learning does not end in the nightmare scenario?

Let me ask Jeff. Jeff, would you take a crack at this?

Jeff Williams:
You know, there have been a few bombastic statements about AI recently. We throw around terms like AI, you know, like the previous question was "where will the AI be, in which chips? Is it in the data center?" I think it's important to define AI a little bit more and decide what we're talking about together.

The fear of the future that you read about, is that somehow machines will take over and start thinking, and I think that that is a long way off, my personal opinion. The very AI brilliance, what we would call brilliance, in the game strategy of beating Go, playing the game Go and winning, you can't take that same learning and play a simple game like tic tac toe and win, so I think we're a long ways off before we can do that kind of human thinking.

As Jensen said, when it comes to specialized deep learning, I think that's going to open up all kinds of opportunities and I think that most of those are not dangerous. I think that the things we need to guard against when you're dealing with large amounts of data which are required to do this type of

learning is really around the privacy and the security of that data.

And so the future I worry about, I don't worry about anytime in the next couple of decades machines starting to think and taking over the world, like you're read about recently. The thing that I think about is the risk to privacy and security and that's one of the reasons that we as an industry have to be really, really smart about what we do. And it's one of the reasons I mentioned earlier that the significance of so much of this being on-device, because I think that provides the greatest level of protection for the individual.

Chairman Chang:
Anybody want to add to that? Yes, Peter.

Peter Wennink:
I'm not any kind of an expert on artificial intelligence. The only thing when – That question is being asked a lot. I'm also on the board of a big pension fund in the Netherlands, and when this question of artificial intelligence and robotics taking over comes up, we need to realize that the demographics in developed countries are tending towards having fewer people who are creating the productivity that we actually need. So we need robotics to take part of that, because I would also say, from a societal point of view, the question is there of inclusion or exclusion. Which people in society are able to follow? That's more of a society question that we need to answer.

And as a matter of fact, the countries that I am close to, Germany and the Netherlands, there's a big, big focus on education, making sure that we get a workforce that can actually deal with these changes and that we don't create a big mass of excluded people. It is also about inclusion, that is something that we need to be cognizant of.

So, workforce is one. We have full employment, effectively, in big parts of Europe. I would welcome some automation, and I would say automated automation.

Chairman Chang:
Thank you. Actually, on this nightmare scenario issue, recently I have been telling the Taiwan press that maybe the first profession that AI will affect may be the medical profession. Because I look at the major hospitals in Taiwan, and

every morning there are hundreds of people registering. In Taiwan, you go to a hospital, you want to see a doctor for whatever reason, unless it's an emergency – if it's an emergency you go to the emergency room. But if it's just an ordinary problem, you have to register.

Every morning before the doctors come to work, hundreds of people will be waiting in the registration room already. They register, and they are called by order of their registration. And most of the ailments are very ordinary – colds, flues, that kind of thing. So what I thought, and what I told the Taiwan press is that maybe AI can really diagnose those minor ailments better than the doctors, and therefore those lines will disappear, and also it will be a relief on the load on the doctors.

But also, at the same time, we may not need so many doctors anymore. Do you have any comment on what I said?

Jeff, you are nodding your head. Do you have any comment?

Jeff Williams:
You know, I think we will always need doctors but I think we will be able to make greater changes in healthcare and empower both the people and the doctors with this. I don't think we're going to put doctors out of work and I don't think you're saying that. But if you think about what doctors do today, most of the doctors we interact with say they only have a few minutes to see each patient, they have very limited data, it's episodic, and there's a big chance to change that in the future.

With the Apple Watch, which has a neural engine on it as well, we're doing more and more. We just announced recently that we're doing a very large study in the United States, and we're going to detect arrhythmias. So the number two cause of death worldwide is stroke, and the largest contributor to stroke is a-fib [atrial fibrillation], and about half of a-fib is asymptomatic, meaning the people who have it don't ever have any symptoms. And with the watch, we've run studies and we've concluded that we can detect this using what you guys broadly call AI, I would say it's closer to machine learning, but doing it on device by looking at the photoplethysmograph. So we envision a world where way less people are dying from stroke because there are interventions if it's detected.

So I agree with Dr. Chang that one of the biggest promises of this is in the world of medicine. And rather than getting rid of all the doctors I think it empowers the doctors to have discussions about how to prevent issues from happening instead of just seeing patients after they do happen.

Chairman Chang:
Actually, I did not say that we would not need so many doctors anymore, but I did say that we would not need so many ordinary doctors anymore. What I meant but did not say, was that we would not need so many mediocre doctors anymore. I did not use the word mediocre, but I used the word ordinary.

Still, I got absolute deafening silence from the medical field anyway.

Jensen Huang:
Morris, one of the largest costs in a hospital is the radiology department, and they're so overworked with the number of scans being taken and the number of images that they have to study that they're outsourcing it to the lowest bidder. And they're doing this in droves.

This is just a perfect example of something that deep learning and AI can help to reduce the burden on and not only will it increase the effectiveness of the diagnosis and early detections, it will also dramatically reduce the cost and burden to the hospitals.

It's been some time, but I think that between the pressure on the hospitals and also the research radiologists, I believe that deep learning is really going to make a difference. There's just a huge movement in the healthcare industry to automate radiology. I think that's a really good direction.

Chairman Chang:
Hock, yes?

Hock Tan:
Just curious, Morris – how do you sue, particularly in the U.S., an AI chip for malpractice?

Chairman Chang:

That is a problem, I think, yes, and I'm a little worried about that. Would you sue the maker of the AI chip? Would you sue the designer of the AI chip? Would you sue the user of the AI chip?

Hock, here is a question from Randy Abrams, who is an analyst with Credit Suisse. His question is: "in an industry, as Hock Tan mentioned, with cloud companies integrating into silicon IP, where are the best opportunities for the chip vendors to create value for these cloud companies and avoid being cannibalized?

Hock Tan:
The ultimate defense to extend the prevailing laws that allow you to do that is really intellectual property. Intellectual property know-how, intellectual property laws. It's really the last line of defense for the chip companies because under my scenario, at the end of the day, what a chip company like Broadcom does can be done by cloud companies drawing on talent out there who are knowledgeable in the specific areas that they are in.

For that matter, OEM companies can do likewise. So the only defense, at the end of the day, is fundamental capabilities in chip design, which relates to IP.

The other defense, which in expanding this question I mentioned, is financials, which has to be built on silicon. Silicon takes a lot of financial resource to make happen. And to be honest, I think at the end of the day the chip company, if they want to stem the reach down the ecosystem stack by cloud guys or OEMs, is to go back really full circle into making chips themselves.

Chairman Chang:
Thank you. Anyone want to add to that? Yes, Steve?

Steve Mollenkopf:
Yes, maybe one thing. I think if you can find problems that are bigger than what the cloud companies can do by themselves, for example, something that creates an entire industry that could just be done by one company – those are interesting opportunities to work on.

Things like privacy, I would say all of the things like 5G that are just large industry fundamentals, they tend to provide opportunities that are bigger than

what any individual company may be interested in and of themselves. I think that if you're just looking for where – hey, I have an advantage because of my size or whatever, it's very difficult to do that with a company that has investment/R&D focus.

So you really want to try to get into something that is changing before anyone else is in it, which is essentially just saying innovation, or innovation at tremendous scale, is really where I think the big opportunities are.

And we talked about a couple of them. I think privacy is one, security, authentication of data – if you don't have those things at an ecosystem level, the industry won't develop, and therefore the cloud companies won't develop services that take advantage of it. I think that's some of the things that the people in this room do well.

Chairman Chang:
Thank you, Steve. Anyone else?

Jensen Huang:
Some of the things that are difficult to replace, and unnecessary to, are platform providers. For example, 5G, at some point, is a platform for the whole industry, ARM is a platform for the whole industry. If you create a platform on which large ecosystems depend on, it's difficult to replace. If you have large scale in doing something, and the scale is so great, and the velocity so high, that no single OEM, even if it's a cloud company, can afford to keep up with it, it's easier and more economic for them to buy yours.

However, if your business is based on products and features, features and benefits, I think those are highly competitive, and it's easy to be integrated, and your features and benefits are no better than somebody else's features and benefits in the final analysis, when it's integrated into a large system. So I think those are difficult to defend. But if you have a platform, or if you have a large scale, or if on the other hand, if you make a trillion little things, those are difficult to replace.

Chairman Chang:
Thank you. Next question is about semiconductor demand. It's from Mr. Masao Asami of Ebara Corporation. The question is: "Where will the most demand

come from for semiconductors in the next 10 years? Self-driving? IOT? Data Centers?

And the second question is “How do we better serve those demands?”

Does anyone want to take a crack at it?

Vince Roche:
I’ll take a crack at it, Morris. I think all those areas are ripe, but I think the one area that I think has been a little dormant in terms of the innovation trajectory over the last 20 years or so is industrial automation. Robotics, the factory automation, process control – there is a lot of anxiety among those customers now to be able to use ubiquitous sensing, to be able to do predictive analysis, predictive maintenance, so on and so forth.

And it’s all about, ultimately, these customers wanting to add more value to their own product portfolios but leverage the technology that they’ve been very, very shy to do for many, many years. So what we’re seeing is a lot more demand for sensing, but I think ultimately, that class of problem, we’ll find it very, very hard to take root if we don’t secure the data right at the sensor. So I think those two things will go hand in glove.

It’s well publicized about autonomous driving, what’s happening there and so forth, but my sense is those initiatives, for example in China, around replacing human capital with machine capital, that’s a very, very real class of opportunity for problems to solve.

Chairman Chang:
Anybody else? Yes, Simon?

Simon Segars:
I think in the near term, IOT is going to drive a lot of volume, and one of the reasons for that is, there exists already today, good enough technology to solve a lot of IOT problems. It needs to get better – we need better nonvolatile memories, we need to be able to integrate all these different technologies, all the sensors, all the digital, all together to make it better, but you can do a lot right now with the technologies that exist. So I think that will drive a lot of volume in the near term.

I think in the long term, I think that semiconductor content is going to go into – I won't say cars but I will say all vehicles to enable truly autonomous driving under all circumstances. There's a long way to go and a lot of R&D required, but it will ultimately consume a ton of silicon. I think that's going to be a big driver in a kind of 10-year plus time frame.

Chairman Chang:
Thank you, Simon. Speaking on behalf of TSMC now, we do have four growth platforms. One is mobile; I think it's going to remain a very strong growth platform for semiconductors, and therefore, for our business.

The second is high-performance computing. The third is IOT, and the third is automotive. These are the four growth platforms that TSMC has right now.

Anyone wants to add anything?

Jensen Huang:
I agree with those four. IOT is super-exciting, because in some way – sensing, IOT, these intelligent machines – we used to have computers, in the future, we're going to have computers that do autonomous things all by themselves.

Most of our computers today we either type into or hold in our hands, but we're going to have intelligent machines everywhere. We're going to have intelligent buildings, we're going to have intelligent factories, factories are going to be AIs, factories are going to run robots that are going to build other robots, and all of these robots are going to be autonomous, entire warehouses will be autonomous, and warehouses will become big AIs, and today, warehouses are non-markets for us. Robots are small markets for us. And in the future, these robots will have dozens of computers inside, not just one ECU that's mimicking the motion that somebody else has programmed. So I think that all of these intelligent machine opportunities made possible by AI are brand new opportunities for us, and it didn't exist before and it's starting from zero, which is the exciting part.

Chairman Chang:
The next question – we really cannot escape Moore's Law in any semiconductor meeting, so the next question is about Moore's Law: "What will

happen after the maximum of Moore's Law? When will that be? What's TSMC strategy or view?"

I don't want to answer that last question, OK? But the first two questions are: "What will happen after the maximum of Moore's Law, and when will that be?"

Peter, do you want to take a crack at that?

Peter Wennink:
When people ask me what's Moore's Law, I say it's an empirical law of economics. It's making sure that the cost per function keeps going down. We either create more functions but then we decrease cost, but it's either/or.

From a patterning or scaling point of view, I think we have a clear vision and roadmap that will bring us to 2030 and further. And beyond that I don't know. It reminds me when I joined the company and Martin van den Brink, my co-President and good friend, he interviewed me and said "Peter, we want you to come." I said "Martin, I have a decent job, why would I join? Because there's Moore's Law, there must be an end to it. So he said "Listen, Moore's Law stops when we run out of ideas, and we have still a lot of ideas left. So another 15 years, if you have the patience, you just join."

I ask him the same question today, he gives me the same answer. So there's this 15 years, the beyond 2030, from a lithography point of view. But like I said in my talk also, I think it is more and more important that we connect the dots.

Complexity is increasing exponentially and it means that we need to look at everything that adds to keeping costs under control, because it is an empirical law of economics. And this is where I think collaboration over the next 10 years needs to take different forms, and I'm not talking about traditional M&A – there is a legislative barrier, I think.

But it is about how you create partnerships where you're willing to share the risks and the rewards. I think that's going to open new venues for the key players in this industry. And I think we will be able to do that.

If we show that we can do that, this is an industry that has massive earnings power, 300 billion dollars of EBIT every year, I think there is massive

knowledge and technology, but we do need to step up the cooperation and collaborative models, and we will make this happen.

Chairman Chang:
Actually, Moore's Law is no longer valid, because Moore's Law has a time element in it. The original version of Moore's Law was doubling of transistor density every 18 to 24 months. The doubling of density is continuing, but not every 18 to 24 months. That's why I say that. The time element has really already disappeared. So that makes the original version of Moore's Law invalid.

But the density doubling is going to continue, and according to the answer that Peter just gave, they have plans all the way to 2030. That's like 13 years more.

Now in my opinion, I think that the economic feasibility of density doubling will be tested first. So before it's technically unfeasible to double the density, the economic viability will be tested first. And that may come before 2030. In fact, TSMC's assumption right now is 2025. Five years more conservative than what Peter just said. 2025 we are expecting that we will continue to work on density doubling for another 8 years.

First of all, the future Chairman and CEO of TSMC are sitting right here. If you disagree with me, please stand up and say that.

As to what will happen after the maximum of Moore's Law – Hock, I think you said something about that. You talked about More than Moore.

Hock Tan:
Well, it will still be a while away, but I think at some point we'll go full circle into a situation where the service providers, clouds, everything else, will start to vertically integrate down to silicon. And at that point I think it's architecture, systems, and software as much as hardware and silicon.

Chairman Chang:
Thank you, Hock. Here's another question. It's from Mr. Taher from OpenSilicon. The question is "as limits of physics come into play, does the panel feel that organic molecules (DNA) could play an important role in increasing the density of data storage in the future?

Does anybody want to answer the question? I don't know enough about organic molecules to answer the question.

Jeff Williams:
My guess, Dr. Chang, is that the question was on DNA. With the level of deep sequencing with all of the data and storage – that's what the genesis of the question was.

Chairman Chang:
I see. Well do you have an answer for the question then?

Jeff Williams:
I was just trying to help (laughs).
You know what's really, really interesting is, when you look at the genome, and you look at all of the sequencing of the genome, the most efficient way to store this data is to store the sample itself, because the way human nature has done it is more efficient that we can project in memory for the next 10 years. You're better off to hold the sample through all the sequencing, versus sequencing and store all the data. The way that nature has done it is better than the technology, at least for the next several years.

Chairman Chang:
Thank you, Jeff.

here's a question about technology again. It says it's for me and for Peter Wennink. Do you see foundry size shrinking through technology or equipment optimization in the future". In other words, can we increase the equipment optimization so that we will need less equipment. I think that is the question.

I have an answer for that – I don't think so. But Peter, you go ahead.

Peter Winnink:
As a matter of fact, we have, over time, sold fewer and fewer machines at higher and higher prices, as you know. But effectively, it comes back to what we said earlier, it has to be able to reduce the cost per transistor or bit or however you want to call it. And I do agree with you, Morris, that Moore's Law will stop sooner through economics than through physics. And that has

continuously been a challenge.

Currently what we have in front of us, and it's just part of the total equation, is lithography, because like what I said earlier, it is a combination of lithography with our place in the industry that provides the material knowledge, that provides the deposition tools, that provides the etchers and everything that we need, and we need to collaborate very closely, very transparently, to really understand what their needs are so that we can work together to solve those technical, but I would also say economical issues.

From a technical point of view, we will shrink. We have our research done, and it's pretty accurate that we are able, the High-NA, to keep shrinking further. That's not the point. The point is can we do it economically, and that's going to determine how big or small the industry is going to be. There's many, many of these elements that need to come together – architectural scaling, device scaling – that all needs to work together.

So it's extremely difficult to give an answer on when does that end, but economics is the big wall. I do agree.

Chairman Chang:
Thank you, Peter.

Now, I have the last question. I would like to close this session in about 5 minutes, so I will save the last question for myself. And the question is principally one for Jeff, I think.

Do you see a lot of AI in mobile devices?

Jeff Williams:
Yes. We talked about some of this earlier, but we do. We think that the frameworks that we've got – we've talked about the neural engines that we've put already in the phone and the watch and given a few examples, but we do view that as a huge piece of the future.

We think these frameworks will allow developers to create apps that will do more in this space. We think that the phone is a major platform for the future. Not exclusively – some pieces will be done in data centers, some pieces will

be done on device, but we're already doing AI in the broader sense of the word, not the machines thinking for themselves sense of AI.

Chairman Chang:
Thank you Jeff. Since Jeff only took a couple minutes, let me ask Vince a question.

Vince, analog devices is really in standard products, isn't that right?

Vince Roche:
It very much depends on the market Morris. Where we build standard products would be in very fragmented applications, number of customers. But clearly as a company we've been moving into more application-specific customer products where the market can justify it and there's a very high concentration of the revenue in a few customers. So it very much depends.

Chairman Chang:
How much of your revenue is in what you call standard products, and how much in the other kind?

Vince Roche:
Well, with the addition of LTC over the past number of months, the amount of standard products in the portfolio has increased. So LT was a very, very standard product play.

All the time we are trying to balance the investment between more standard product, more open market type things, versus more customization. But clearly the trend at ADI has been more customization over the past 10 years.

Chairman Chang:
Let me get to my question then. One usually associates low profit margin with standard products. How does Analog Devices with so many standard products?

Vince Roche:
It's a great industry secret. (laughs)

I think it depends on the market. Not every market produces the same

economics for ADI. So it very much depends on the type of product and the type of market in which you're playing. Clearly the more fragmented markets offer more attractive margins structures and the more concentrated markets...you know the buyer has some very, very strong power, so the seller has to do a lot more bargaining.

Chairman Chang:
Thank you. And now we end this panel session. I want to thank all the panelists. Thank you for participating. You have been a great panel, thank you.